arXiv:1606.01548v1 [math.OC] 5 Jun 2016

EXACT RECOVERY WITH SYMMETRIES FOR PROCRUSTES MATCHING

NADAV DYM AND YARON LIPMAN

Abstract. The Procrustes matching (PM) problem is the problem of finding the optimal rigid motion and labeling of two point sets so that they are as close as possible. Both rigid and non-rigid shape matching problems can be formulated as PM problems. Recently [14] presented a novel convex semi-definite programming relaxation (PM-SDP) for PM which achieves state of the art results on common shape matching benchmarks.

In this paper we analyze the successfulness of PM-SDP in solving PM problems without noise (Exact PM problems). We begin by showing Exact PM to be computationally equivalent to the graph isomorphism problem. We demonstrate some natural theoretical properties of the relaxation, and use these properties together with the moment interpretation of [12] to show that for exact PM problems and for (generic) input shapes which are asymmetric or bilaterally symmetric, the relaxation returns a correct solution of PM.

For symmetric shapes, PM has multiple solutions. The non-convex set of optimal solutions of PM is strictly contained in the convex set of optimal solutions of PM-SDP, so that 'most' solutions of PM-SDP will not be solutions of PM. We deal with this by showing the solution set of PM to be the extreme points of the solution set of PM-SDP, and suggesting a random algorithm which returns a solution of PM with probability one, and returns all solutions of PM with equal probability. We also show these results can be extended to the almost-exact case. To the best of our knowledge, our work is the first to achieve exact recovery in the presence of multiple solutions.

1. Introduction. Shape comparison is a central task in many fields such as computer graphics, computer vision, medical imaging and biology in general. The input of the problem is a pair of shapes, often represented by d-dimensional point clouds $P \in \mathbb{R}^{d\times n}$, $Q \in \mathbb{R}^{d\times n}$, where a point is represented by the column of a matrix. The goal is to determine how similar the shapes are to each other, typically by computing a distance between shapes which is invariant to a prescribed class of shape preserving deformations.

In rigid shape matching problems, shape preserving deformations are rigid motions, as well as permutations of the point cloud which change the order the points are given but not the shape defined by them. A distance between shapes for rigid problems can be defined by finding a linear isometry $R \in \mathcal{O}(d)$, a translation $t \in \mathbb{R}^d$ and a permutation $X \in \Pi_n$ minimizing the Euclidian distance between the point clouds. Centralizing the point clouds causes the optimal translation to be $t = 0$, and we arrive at the following optimization problem:

$$\mathrm{d}^2(P,Q) = \min_{X,R} \|RP - QX\|_F^2 \tag{1a}$$

$$\text{s.t.} \quad X \in \Pi_n \tag{1b}$$

$$R \in \mathcal{O}(d) \tag{1c}$$

We will refer to this optimization problem as the Procrustes matching (PM) problem. We note that both the minimal value and the minimizers are of practical importance. The minimal value is a measure of the similarity of the shapes, while the minimizers R, X enable correct alignment and labeling of the shapes.

For non-rigid shapes, shape preserving deformations also include intrinsic non-rigid isometries. In computer graphics, a popular strategy for tackling the non-rigid problem is embedding the shapes in a higher dimensional space, where non-rigid isometries of the shapes are approximated by linear isometries in the high dimensional

FIG. 1. *Examples of non-rigid matching results obtained by PM-SDP. Results are visualized by transforming a color map on the source shape (left) to the target shape (right) according to the correspondence found by PM-SDP.*

space (*e.g.*, [15]) . As a result the non-rigid problem can also be formulated as PM, though now d is high dimensional as opposed to the rigid case where typically $d = 3$.

PM is a non-convex optimization problem and global minimization is difficult. In fact we show that even the subproblem of exact PM is difficult. Exact PM refers to the problem of deciding whether $\mathrm{d}(P, Q) = 0$ or not, and the related problem of finding minimizers (R, X) to (1) when the Procrustes distance is zero. We show this by proving

THEOREM 1. *There is a polynomial reduction from exact GM to exact PM, and vice versa.*

Exact GM is a well studied problem which is not yet known to be either polynomial or NP-complete, although recently it has been shown to be solvable in quasi-polynomial time [3]. In [14] we presented a method for approximating the solution of PM by formulating a semi-definite convex relaxation which we name PM-SDP. The relaxation is constructed using standard methods for semi-definite relaxations of quadratic problems. Using results on semi-definite matrix completion, the relaxation is reduced to a considerably more efficient SDP relaxation which is equivalent to the original relaxation. The usefulness of PM-SDP for rigid and non-rigid matching problems was demonstrated as well. Some examples of non-rigid shape matching using PM-SDP are presented in Figure 1.

Our main goal in this paper is to theoretically justify the successfulness of PM-SDP on exact or near exact PM problems. This successfulness was demonstrated empirically by the non-rigid results in [14] and by experiments presented there which

showed that for exact and almost-exact problems the correct solution of PM was retrieved. We show that while solving exact PM is computationally hard in general, under assumptions which are typically valid for shape matching problems, PM-SDP is guaranteed to correctly solve exact and near-exact PM problems. In particular our assumptions are valid for 'most' asymmetric or bilaterally symmetric shapes. The latter class includes many important shape matching instances such as humans and other animals; for example, all the shapes in Figure 1 are bilaterally symmetric. For simplicity of the exposition we currently refer to P which satisfy our assumptions as 'generic', and explain our assumptions afterwards.

Apart from the exactness results, we also show some natural properties of the relaxation. We show invariance of the algorithm to orthogonal transformation and relabeling of the input, that the R, X coordinates of the solution are always in the convex hull of $\mathcal{O}(d) \times \Pi_n$, and that the objective value $\underline{\mathrm{d}}(P, Q)$ of PM-SDP is always non-negative, and is never larger that $\mathrm{d}(P, Q)$.

1.1. Main results. From the properties of the relaxation we just mentioned it follows that for exact problems

$$0 \leq \underline{\mathrm{d}}(P, Q) \leq \mathrm{d}(P, Q) = 0$$

so that PM-SDP attains the correct objective value for exact problems. Our goal is to show that PM-SDP attains the correct minimizers as well.

For exact problems, we will refer to minimizers of PM as *exact solutions* and will call minimizers of PM-SDP *convex exact solutions*. Note that the set of *convex exact solutions* is a convex set containing the set of exact solutions. For generic asymmetric point-clouds we show that these two sets are equal:

THEOREM 2 (asymmetric version). *Assume P is generic and asymmetric, and $\mathrm{d}(P, Q) = 0$. Then PM-SDP has a unique convex exact solution, which coincides with the unique exact solution of PM.*

For symmetric P,Q, there are several exact solutions. Each such solution is also a convex exact solution. In fact, all convex combinations of these solutions will be convex exact solutions as well, so that in the symmetric case convex exact solutions will not necessarily be exact solutions. To deal with this problem we will need some notation:

We denote the set of exact solutions by $\mathrm{ISO}(P, Q)$ and the projection of this set onto the R coordinate by $\mathrm{ISO}_R(P, Q)$. We refer to members of this set as exact orthogonal solutions. Similarly we denote the set of convex exact solutions by $\mathcal{N}(P, Q)$ and its projection onto the R coordinate by $\mathcal{N}_R(P, Q)$. We refer to the members of this set as convex exact orthogonal solutions. We prove that the extreme points of the set of convex exact orthogonal solutions is exactly the convex hull of the set of exact orthogonal solutions:

THEOREM 2 (symmetric version). *Assume P is generic and $\mathrm{d}(P, Q) = 0$. Then*

$$\mathrm{ISO}_R = \mathrm{ext}(\mathcal{N}_R)$$

where $\mathrm{ext}(\mathcal{N}_R)$ denotes the extreme points of $\mathcal{N}_R$.

Under the condition of the theorem, we obtain a semi-definite characterization of the convex hull of the optimal set of PM, as the intersection of the feasible set of PM-SDP with the hyperplane defined by constraining the objective to be zero. We can then obtain extreme points (=exact orthogonal solutions) by a random algorithm

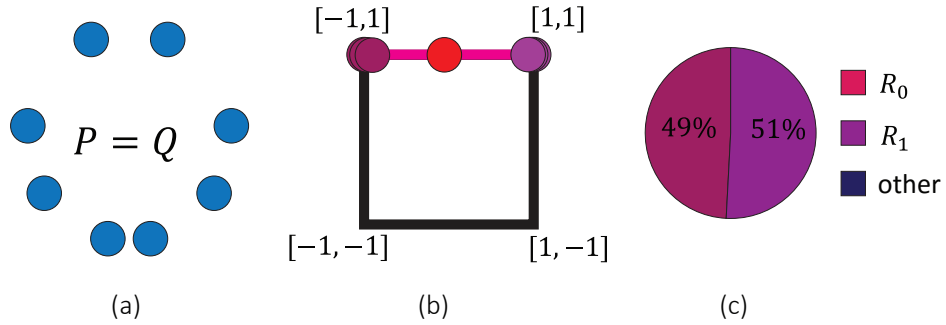

FIG. 2. *An experiment illustrating our main result. We apply PM-SDP to bilaterally symmetric point clouds $P = Q$ seen in (a). PM has two solutions $(R_i, X_i), i = 0, 1$, where R_i are diagonal matrices with values $[1, 1]$ or $[-1, 1]$ on the diagonal. The solution of PM-SDP (the red dot in (b)) is a convex combination of R_0 and R_1 as predicted by Theorem 2. By applying a random projection the correct solutions of PM (purple dots in (b)) are obtained with equal probability as seen in the pie chart in (c). This is in accordance with Theorem 3.*

which maximizes random linear energies over this set. This strategy for obtaining extreme points was suggested by Barvinok in [5]. We prove that if the linear energies are selected according to the uniform distribution on the unit sphere, then

THEOREM 3. *The random algorithm returns an exact solution with probability one. Moreover the probability distribution induced on the exact solutions is uniform.*

We also show the random algorithm can be adapted to the almost exact case, and maintain similar theoretical guarantees.

An experiment illustrating our main results is shown in Figure 2.

Related work. A popular strategy for local minimization of PM is the ICP algorithm [6, 16] which iteratively solves for R and then X, while holding the other coordinate fixed. ICP enjoys excellent scalability, but requires good initialization to avoid local minima. Another popular strategy is RANSAC [8] which repeatedly randomly selects $d+1$ points from each shape and finds the optimal rigid motion between them. The global rigid motion is then found or approximated by applying a voting scheme to a sufficient number of trials. For high dimensional data this method is problematic since its complexity is exponential in d.

[11] propose a natural convex relaxation to PM by optimizing over the convex hull of the feasible set $\mathcal{O}(d) \times \Pi_n$. This relaxation is faster than PM-SDP but is provably less accurate (see [10] for a similar argument for the graph matching problem). A drawback of this relaxation is that while exact recovery is guaranteed for certain asymmetric shapes, it is also known to fail for centralized point clouds since in this case $R = 0, X = \frac{1}{n}\mathbf{1}\mathbf{1}^T$ is always a feasible zero-objective solution.

There are many examples of convex relaxations of non-convex problems for which exact recovery results are available, *e.g.*, [17, 9, 1], often under certain assumptions on the data. The papers mentioned above all address situations in which the non-convex problem has a unique solution and show that the relaxation has a unique solution as well. To the best of our knowledge our treatment here is the first proof of exact recovery for a relaxation of symmetric problems which possess multiple solutions.

1.2. Conditions. We now explain our assumptions on the point cloud P. A weak assumption we use is that P does not include the same point twice. Our main assumption is that the spectrum of the matrix PP^T is simple (all eigenvalues have

unit multiplicity). This assumption implies that exact orthogonal symmetries are compositions of reflections with respect to the principal axes of P. In other words, in the principal axes basis an orthogonal symmetry is diagonal, with diagonal entries in $\{-1,1\}$. In particular all symmetries are bilateral ($R^2 = I_d$). We say the R is a member of $\{-1,1\}^d$ if it is such a diagonal matrix.

The simple spectrum assumption and the weaker bounded eigenvalue multiplicity assumption are common assumptions in the exact graph matching literature [4, 1]. It is known that in this case exact GM can be solved by an algorithm which is polynomial in the input size but exponential in the eigenvalue multiplicity. This suggests that exact recovery for PM-SDP should fail at some level of eigenvalue multiplicity. Indeed we provide examples where such failures occur even for multiplicity two.

Our third and final assumption is the faithfulness assumption. We say a point P_j is *faithful*, if for any $R \in \{-1,1\}^d$ for which RP_j is a point of P, R is an orthogonal symmetry. We say that P is faithful if it contains at least one faithful point.

An illustration of the faithfulness condition is given in the inset. The image on the left of the inset shows a shape from the Scape dataset [2] which when sampled has a faithful point marked in green. Applying the reflectional symmetry $R \in \mathrm{ISO}_R(P,P)$ to the point in green gives another green point on the shape, while applying $R \in \{-1,1\}^2$ which aren't symmetries gives points which aren't on the shape (dotted, light green). The image on the right shows a synthetic example of a point-cloud where all points are unfaithful, since the point-cloud is asymmetric but each point can be reflected to a different point in the point cloud.

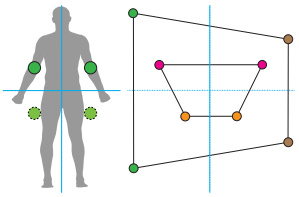

Our experiments show that PM-SDP is able to correctly solve the unfaithful exact PM problem displayed in the inset, as well as similar randomly generated problems. Thus, in contrast to the simple spectrum condition which is tight, we do not currently know whether the faithfulness condition can be removed or replaced by a weaker condition.

1.3. Paper organization. We begin by presenting our SDP relaxation for PM in Section 2. We prove the theoretical properties of the relaxation in Section 3 and use them in Section 4 to prove Theorem 2. In Section 5 we describe the random algorithm for obtaining extreme points and prove Theorem 3. We then explain how this result can be extended to accommodate near exact problems. In Section 6 we present some experiments which illustrate our results, and examine the behavior of PM-SDP for exact problems which don't meet the conditions of our analysis. Finally we show the equivalence of exact graph matching and exact PM in Section 7.

1.4. Notation. We denote the set $\{1,\dots,n\}$ by $[n]$.

We denote polynomials of degree $\leq r$ in $x = (x_1,\dots,x_N)$ by $\mathcal{P}_r(x)$.

We denote the feasible set of (1) by $G = \mathcal{O}(d) \times \Pi_n$.

We denote the j-th column of a matrix $X \in \mathbb{R}^{n\times n}$ by $X_j \in \mathbb{R}^{n\times 1}$, and the i-th row by $X_{i\star} \in \mathbb{R}^{1\times n}$. Expressions such as X_j^T should be interpreted as $(X_j)^T$ (as opposed to $(X^T)_j$).

We denote by $\mathbf{1}$ the vector $\mathbf{1} = (1,1,\dots,1)^T \in \mathbb{R}^{n\times 1}$. All vectors are column vectors unless stated otherwise.

Real symmetric $n \times n$ matrices are denoted by $\mathcal{S}(n)$. We use $A \succeq 0$ to denote positive semi-definite matrices.

2. Convex relaxation. PM can be formulated using quadratic polynomials in the entries of R, X. For generality of the exposition, as well as for notational convenience, we will now review a common strategy for semi-definite relaxation of quadratic optimization problems (see [13] for a review) and will then state the formulation of the PM-SDP relaxation from [14]. We present the relaxation using the terminology of [12], which we find helpful for noting properties of the relaxation which are otherwise less apparent.

2.1. SDP relaxations for quadratic optimization problems. We consider quadratic optimization problems of the form

$$\min_{x\in\mathbb{R}^N} \quad f_0(x) \tag{2a}$$

$$\text{s.t.} \quad f_s(x) = 0, \quad s = 1,\dots,S \tag{2b}$$

where f_s, $s = 0,\dots,S$ are quadratic polynomials. We denote the set of points x satisfying (2b) by K. We will assume K is compact. This non-convex optimization problem can be rephrased as an equivalent convex problem over the infinite dimensional space $\mathbb{P}(K)$ of probability measures supported in K:

$$\min_{\mu\in\mathbb{P}(K)} \mu[f_0] = \int f_0 d\mu \tag{3}$$

The equivalence is in the sense that the optimal values of both problems are equal, and each solution x of (2) defines a solution δ_x of (3). Solutions of (3) can also be (possibly continuous) convex combinations of such solutions δ_x. We use this equivalence to construct a finite dimensional convex relaxation of (2). Instead of considering $\mathbb{P}(K)$ we will consider a superset $\mathcal{M}$ consisting of (not necessarily positive) linear functionals on $C(K)$. A member $\mu \in \mathcal{M}$ must fulfill the following conditions:

1. $\mu[1] = 1$
2. For all s, $\mu[f_s] = 0$, for $s = 1,\dots,S$.
3. For all $p \in \mathcal{P}_1(x)$, $\mu[p^2] \geq 0$. If μ satisfies this property we write $\mu \succeq 0$.

Note that indeed $\mathbb{P}(K) \subset \mathcal{M}$ since every probability measure on K satisfies these conditions; in fact it satisfies the last condition for every non-negative integrable function. Applying this relaxation to (3) gives

$$\min_{\mu\in\mathcal{M}} \quad \mu[f_0] \tag{4a}$$

$$\text{s.t.} \quad \mu[1] = 1 \tag{4b}$$

$$\mu[f_s] = 0, \quad s = 1\dots S \tag{4c}$$

$$\mu \succeq 0 \tag{4d}$$

To see this is indeed a finite dimensional semi-definite program note that in (4) μ is only applied to polynomials in $\mathcal{P}_2(x)$ and therefore the unknowns of the problem are

$$\left(\mu[1], \mu[x_1], \dots, \mu[x_N], \mu[x_1^2], \mu[x_1x_2], \dots, \mu[x_N^2]\right) \tag{5}$$

Equations (4a)-(4c) are all linear in the unknowns, while (4d) is a semi-definite constraint as explained next. For $p, q \in \mathcal{P}_1(x)$

$$p(x) = \mathbf{p}^T \begin{pmatrix} 1 \\ x \end{pmatrix} \quad , \quad q(x) = \mathbf{q}^T \begin{pmatrix} 1 \\ x \end{pmatrix}$$

where $\mathbf{p}, \mathbf{q}$ are the coefficient vectors of p, q. Then

$$\mu[pq] = \mu\left[\mathbf{p}^T \begin{pmatrix} 1 \\ x \end{pmatrix} \begin{pmatrix} 1 & x^T \end{pmatrix} \mathbf{q}\right] = \mathbf{p}^T \mu \begin{bmatrix} 1 & x^T \\ x & xx^T \end{bmatrix} \mathbf{q}$$

Where for a function $F : \mathbb{R}^N \to \mathbb{R}^{n \times k}$ such that $F_{ij}(x) \in \mathcal{P}_2(x)$ we define $\mu[F]$ by applying μ to each coordinate, *i.e.*, $\mu[F]_{ij} = \mu[F_{ij}]$. It follows from (2.1) that

$$\mu \succeq 0 \text{ if and only if } \mu \begin{bmatrix} 1 & x^T \\ x & xx^T \end{bmatrix} \succeq 0$$

We denote the set of feasible solutions of (4) by $\mathcal{M}_F$.

2.2. Properties. We present some simple but important consequences of the discussion above:

PROPOSITION 4. *Assume $\mu \in \mathcal{M}_F$ and $p \in \mathcal{P}_1(x)$,*

1. *If $\mu[p^2] = 0$ then $\mu[pg] = 0$ for all $g \in \mathcal{P}_1(x)$.*
2. *$\mu[p^2] \geq \mu[p]^2$.*

Proof. For any positive c,

$$\mu\left[(c^{-1}p \pm cg)^2\right] \geq 0$$

which implies that

$$2|\mu[pg]| \leq c^{-2}\mu[p^2] + c^2\mu[g^2]$$

The first claim follows by taking $c \to 0$. The second claim follows by choosing $c = 1$ and taking g to be the constant function $g = \sqrt{\mu[p^2]}$. ∎

An immediate consequence is that if $p_1, \ldots, p_k \in \mathcal{P}_1(x)$ all satisfy $\mu[p_i^2] = 0$, then $\mu[f] = 0$ for all members f of the vector space

$$\langle p_1, \ldots, p_k \rangle_2 = \{f | f = \sum_{i=1}^{k} p_i g_i, \quad g_i \in \mathcal{P}_1(x)\}$$

2.3. Convex relaxation of the Procrustes problem. The PM problem is equivalent to the quadratic optimization problem:

$$\min_{X,R} \quad \|RP - QX\|_F^2 \tag{6a}$$

$$\text{s.t.} \quad X\mathbf{1} = \mathbf{1} \quad , \quad \mathbf{1}^T X = \mathbf{1}^T \tag{6b}$$

$$X_j X_j^T = \operatorname{diag}(X_j), \quad j = 1 \ldots n \tag{6c}$$

$$RR^T = I_d \quad , \quad R^T R = I_d \tag{6d}$$

where for a vector v, $\operatorname{diag}(v)$ is the diagonal matrix whose diagonal entries are the entries of v. We note that (6c) follows from the fact that the entries of each column X_j are in $\{0, 1\}$, and each column has only one non-zero entry. The equivalence of PM with (6) is explained in full detail in [14].

Applying the semi-definite relaxation described in the previous section to the

quadratic formulation of PM we obtain our formulation for PM-SDP:

$$\underline{\mathrm{d}}^2(P,Q) = \min_{\mu} \quad \mu\left[\|RP - QX\|_F^2\right] \tag{7a}$$

$$\text{s.t.} \quad \mu[1] = 1 \tag{7b}$$

$$\mu[X]\mathbf{1} = \mathbf{1} \quad , \quad \mathbf{1}^T\mu[X] = \mathbf{1}^T \tag{7c}$$

$$\mu[X_j X_j^T] = \mathrm{diag}\left(\mu[X_j]\right), \quad j = 1 \dots n \tag{7d}$$

$$\mu[RR^T] = I_d \quad , \quad \mu[R^T R] = I_d \tag{7e}$$

$$\mu \succeq 0 \tag{7f}$$

An important property of PM-SDP (which is not needed for our discussion here) is that (7f) can be replaced with the semi-definite constraints

$$\mu\begin{bmatrix} 1 & x_j^T \\ x_j & x_j x_j^T \end{bmatrix} \succeq 0, \quad j = 1, \dots, n \tag{8}$$

to obtain an *equivalent* relaxation. Here x_j is a vector containing the entries of R and the column X_j. In applications often $n >> d$, and in this case the $O(n + d^2)$ sized matrices involved in (8) are considerably smaller than the $O(n^2 + d^2)$ matrix in (7f). As a result the equivalent relaxation is considerably more efficient than (7). For an explanation of the equivalence see [14].

3. Properties of relaxation. We present some natural properties of the relaxation, which will be helpful for the proof of our main theorem. We begin by presenting two consequences from our discussion above.

PROPOSITION 5. *For $\mu \in \mathcal{M}_F$, the objective (7a) satisfies*

$$\mu\left[\|RP - QX\|_F^2\right] \geq \|\mu[R]P - Q\mu[X]\|_F^2$$

In particular $\underline{\mathrm{d}}$ is bounded from below by zero, and if $\underline{\mathrm{d}}(P,Q) = 0$ then

$$\mu[R]P = Q\mu[X]$$

Proof of proposition 5. We note that $\|RP - QX\|_F^2$ can be rewritten as $\sum_{i,j} p_{ij}^2$ for $p_{ij} = (RP - QX)_{ij} \in \mathcal{P}_1(x)$. Using Proposition 4 we obtain

$$\mu\left[\|RP - QX\|_F^2\right] = \sum_{i,j} \mu\left(p_{ij}^2\right) \geq \sum_{i,j} \mu\left(p_{ij}\right)^2 = \|\mu[R]P - Q\mu[X]\|_F^2 \geq 0$$

∎

PROPOSITION 6. *$\mu[R], \mu[X]$ are in the convex hull of the orthogonal transformations and permutations, respectively.*

Proof of proposition 6 . We begins by showing that $\mu[X]$ is in the convex hull of Π_n, *i.e.*, that $\mu[X]$ is doubly stochastic. The rows and columns of $\mu[X]$ are constrained to sum to one in (7c), and each coordinate of $\mu[X]$ must be non-negative since by (7d)

$$\mu[X_{ij}] = \mu[X_{ij}^2] \geq 0$$

The convex hull of orthogonal matrices are matrices whose 2-norm is not larger than one. For arbitrary $v \in \mathbb{R}^d$ using Proposition 4 and (7e) we have

$$\|\mu[R]v\|_2^2 \leq \mu\left[\|Rv\|_2^2\right] = \mu\left[v^T R^T R v\right] = v^T \mu[R^T R] v = \|v\|_2^2$$

and therefore $\|\mu[R]\|_2 \leq 1$. ∎

3.1. Compactness. We show that $\mathcal{M}_F$ is compact. To show that all coordinates of a semi-definite matrix are bounded, it is sufficient to show that its trace is bounded. Therefore for general quadratic relaxations it is sufficient to show that $\sum_i \mu[x_i^2]$ is bounded. In our case x consists of the coordinates of R, X, and so the compactness of $\mathcal{M}_F$ follows from

$$\sum_{ij} \mu[X_{ij}^2] + \sum_{k,\ell} \mu[R_{k\ell}^2] = \sum_{ij} \mu[X_{ij}] + \sum_{k} \mu[(RR^T)_{kk}] = n + d$$

3.2. Invariance to coordinate change and reordering. The Procrustes distance $\mathrm{d}(P, Q)$ is invariant to orthogonal change of coordinates and reordering of the points, that is given $(R_0, X_0), (R_1, X_1) \in G$ we have

$$\mathrm{d}(P, Q) = \mathrm{d}(\hat{P}, \hat{Q}) \tag{9}$$

where $\hat{P} = R_0 P X_0$ and $\hat{Q} = R_1 Q X_1$. We now show that our convex approximation $\underline{\mathrm{d}}$ satisfies (9) as well.

We define

$$(R_0, R_1, X_0, X_1)_*(\mu) = \hat{\mu}$$

where

$$\hat{\mu}[p(R, X)] = \mu[p(R_1 R R_0^T, X_1^T X X_0)]$$

It can be verified that if $\mu \in \mathcal{M}_F$ then $\hat{\mu} \in \mathcal{M}_F$ as well. Also note that $(R_0, R_1, X_0, X_1)_*$ is a linear map whose inverse is $(R_0^T, R_1^T, X_0^T, X_1^T)_*$. Finally, note that

$$\left\| R\hat{P} - \hat{Q}X \right\|_F^2 = \left\| (R_1^T R R_0) P - Q(X_1 X X_0^T) \right\|_F^2$$

which implies that

$$\hat{\mu}\left[\left\| R\hat{P} - \hat{Q}X \right\|_F^2 \right] = \mu\left[\|RP - QX\|_F^2 \right]$$

It follows that (9) holds, and there is a linear isomorphism taking the minimizers of PM-SDP(P, Q) to the minimizers of PM-SDP$(\hat{P}, \hat{Q})$.

4. Exact recovery. In this section we prove our main theorem (Theorem 2). We are given P, Q which satisfy the conditions of the theorem. Our goal is to show that the set of exact orthogonal solutions $\mathrm{ISO}_R(P, Q)$ is equal to the extreme points of the set of convex exact orthogonal solutions $\mathcal{N}_R(P, Q)$. In fact it is sufficient to prove that the convex hull of $\mathrm{ISO}_R(P, Q)$ is equal to $\mathcal{N}_R(P, Q)$. This is because the extreme points of $\operatorname{conv} \mathrm{ISO}_R(P, Q)$ are precisely $\mathrm{ISO}_R(P, Q)$, which in turn follows from the fact that all orthogonal matrices are extreme points of $\operatorname{conv}(\mathcal{O}(d))$.

The inclusion $\operatorname{conv} \mathrm{ISO}_R(P, Q) \subseteq \mathcal{N}_R(P, Q)$ follows immediately from the fact that each exact orthogonal solution is also a convex exact orthogonal solution.

The proof of the opposite inclusion $\mathcal{N}_R(P, Q) \subseteq \operatorname{conv} \mathrm{ISO}_R(P, Q)$ is more involved. The remainder of this section is devoted to proving this inclusion.

We first note that using the invariance of PM-SDP to applying permutations and orthogonal transformations we can assume w.l.o.g. that $P = Q$ since by assumption $\mathrm{d}(P, Q) = 0$. Additionally we can assume that PP^T is a diagonal matrix denoted by D. To see this we use the spectral decomposition

$$UPP^TU^T = D$$

and note that $\hat{P} = UP$ satisfies $\hat{P}\hat{P}^T = D$.

By multiplying the equation $RP = PX$ with its transpose it can be seen that all exact solutions $(R, X) \in \mathrm{ISO}(P, P)$ satisfy $RD = DR$. In the simple spectrum case this implies that R is diagonal and $R_{jj} \in \{-1, 1\}$. Thus the symmetry group of P can be identified with a subgroup of $\{-1, 1\}^d$. The proof of the theorem is based on the following generalization of these properties of exact solutions to exact convex solutions:

LEMMA 7. *Assume $\mu \in \mathcal{N}$. Then*

1. $\mu\left[\|RD - DR\|_F^2\right] = 0$.
2. $\mu\left[R_{ij}^2\right] = \delta_{ij}$.
3. *$\mu[R]$ is diagonal, and $\mu[R]_{jj} \in [-1, 1]$.*
4. *If $\mu[X_{\ell j}] > 0$ then there is some $R \in \{-1, 1\}^d$ such that $RP_j = P_\ell$.*

Proof. For given $\mu \in \mathcal{N}$ we have

$$\mu\left[\|RD - DR\|_F^2\right] = \mu\left[\mathrm{tr}\left(DR^T RD\right)\right] - 2\mu\left[\mathrm{tr}\left(RDR^T D\right)\right] + \mu\left[\mathrm{tr}\left(DRR^T D\right)\right]$$

We show the expression above vanishes by showing that all three summands in the last expression are equal to $\|D\|_F^2$. For the first and third summands this follows immediately from the linearity of μ and the constraint (7e). For the second summand, recall that $D = PP^T$ and note that

$$\mathrm{tr}\left(RPP^T R^T PP^T\right) - \mathrm{tr}\left(PXX^T P^T PP^T\right) \in \langle RP - PX\rangle_2$$

Therefore

$$\mu\left[\mathrm{tr}\left(RDR^T D^T\right)\right] = \mathrm{tr}\left(P\mu[XX^T]P^T PP^T\right) = \|D\|_F^2$$

where for the last equality on the right we use the fact that

$$\mu[XX^T] = \sum_j \mu[X_j X_j^T] = \sum_j \mathrm{diag}\left(\mu[X_j]\right) = I_n$$

This proves the first claim. We now have that

$$0 = \mu\left[(RD - DR)_{ij}^2\right] = (D_{jj} - D_{ii})^2 \mu[R_{ij}^2]$$

Since the diagonal elements of D are distinct by the simple spectrum assumption, we see that for $i \neq j$ we have $\mu[R_{ij}^2] = 0$ while when $i = j$ we have

$$\mu[R_{ii}^2] = \sum_j \mu[R_{ij}^2] = \mu[(RR^T)_{ii}] = 1$$

This proves the second claim. Using Proposition 4 it follows immediately that $\mu[R]$ is diagonal, and $\mu[R]_{jj} \in [-1, 1]$.

We now prove our last claim. We assume $\mu[X_{\ell,j}] > 0$, and our goal is to show that for any $i \in [d]$, $|P_{i\ell}| = |P_{ij}|$. We fix some i and define a vector v by

$$v_k = P_{ik}^2$$

Our goal is to show that $v_\ell = v_j$. Note that for all k,

$$v_k = P_{ik}^2 \mu[R_{ii}^2] \overset{(a)}{=} \mu[(RP)_{ik}^2] \overset{(b)}{=} \mu[(PX)_{ik}^2] \overset{(c)}{=} v\mu[X_k]$$

where:

(a) follows from the fact that $(RP)^2_{ik} - P^2_{ik}R^2_{ii}$ is a member of $\langle R_{ij} |\ j \neq i\rangle_2$.

(b) follows from the fact that

$$(RP)^2_{ik} - (PX)^2_{ik} = (RP - PX)_{ik}(RP + PX)_{ik} \in \langle (RP - PX)_{ik} \rangle_2$$

and (c) follows from

$$\mu[(PX)^2_{ik}] = \mu[P_{i\star}X_k X_k^T P_{i\star}^T] = P_{i\star}\mathrm{diag}\left(\mu[X_k]\right) P_{i\star}^T = v^T \mu[X_k] \qquad \square$$

As (4) holds for all k, it follows that $v^T = v^T\mu[X]$. Since $\mu[X]$ is doubly stochastic it can represented as a convex combination of permutation matrices $X(s)$. Accordingly v^T is a convex combination of the vectors $v^T X(s)$ whose norm is no larger than $||v||$, and this implies that $v^T X(s) = v^T$ for all s. There must be some s for which $X_{\ell j}(s) = 1$ and therefore

$$v_j = (v^T X(s))_j = v_\ell$$

We now explain how the theorem is proved from the lemma. Let us assume w.l.o.g. that P_1 is the faithful column of P. Let $\mathcal{J}$ be the collection of indices ℓ such that $\mu[X_{\ell 1}]$ is strictly positive. By the last part of the lemma, for each $\ell \in \mathcal{J}$ there is some $R(\ell) \in \{-1, 1\}^d$ such that $R(\ell)P_1 = P_\ell$. Moreover due to faithfulness $R(\ell)$ is an exact orthogonal solution.

Now note that for all $i \in [d]$,

$$\mu[R_{ii}]P_{i1} = \mu[(RP)_{i1}] = \mu[(PX)_{i1}] = \left(\sum_{\ell \in \mathcal{J}} \mu[X_{\ell 1}] R_{ii}(\ell)\right) P_{i1}$$

Note that if all coordinates of P_1 are non-zero then for each i we can cancel out P_{i1}, so that $\mu[R]$ is a convex combination of exact orthogonal solutions as required:

$$\mu[R] = \sum_{\ell \in \mathcal{J}} \mu[X_{\ell 1}] R(\ell) \tag{10}$$

For the general case we define $\hat{R}(\ell)$ to be the diagonal matrix with diagonal elements

$$\hat{R}(\ell)_{ii} = \begin{cases} R(\ell)_{ii} & \text{if } P_{i1} \neq 0 \\ \mu[R]_{ii} & \text{if } P_{i1} = 0 \end{cases}$$

We note that (10) holds when $R(\ell)$ is replaced with $\hat{R}(\ell)$. Thus it is sufficient to show that each $\hat{R}(\ell)$ is a convex combination of exact orthogonal solutions.

Since the diagonal elements of $\mu[R]$ are in $[-1, 1]$, each $\hat{R}(\ell)$ can be written as a convex combination of matrices $R(\ell, \alpha) \in \{-1, 1\}^d$ satisfying the condition

$$R(\ell, \alpha)_{ii} = \hat{R}(\ell)_{ii} = R(\ell)_{ii} \text{ if } P_{i1} \neq 0$$

Note that $R(\ell, \alpha)P_1 = R(\ell)P_1 = P_\ell$ and so all matrices $R(\ell, \alpha)$ are exact orthogonal solutions due to faithfulness. This concludes the proof of the symmetric version of the theorem.

In the asymmetric case, what we showed is that PM-SDP returns the unique exact orthogonal solution of PM. It remains only to show that the X coordinate obtained from PM-SDP agrees with the exact solution as well. This follows from the following lemma

LEMMA 8. *Assume P, Q satisfy $\mathrm{d}(P,Q) = 0$ and Q has no repeated points. Let $(R(0), X(0)) \in G$ be an exact solution for PM. Then $X = X(0)$ is the unique doubly stochastic solution to the equation*

$$R(0)P = QX \tag{11}$$

Proof. Note that since the columns of Q are pairwise disjoint $X(0)$ is the unique solution of (11) over the set of permutation. It remains to show this is the case over the set of doubly stochastic matrices as well.

Note that for any doubly stochastic X,

$$0 \leq \|R(0)P - QX\|_F^2 = \|R(0)P\|_F^2 - 2\mathrm{tr}R(0)PX^TQ^T + \|QX\|_F^2$$

and therefore

$$\mathrm{tr}R(0)PX^TQ^T \leq \frac{1}{2}\left(\|R(0)P\|_F^2 + \|QX\|_F^2\right) \leq \frac{1}{2}\left(\|P\|_F^2 + \|Q\|_F^2\right) = \|P\|_F^2 \qquad \square$$

Therefore, the linear functional $\mathrm{tr}R(0)PX^TQ^T$ is maximized by $X = X(0)$ and for any other permutation it is suboptimal. It is therefore suboptimal also for any other matrix in the convex-hull of the permutations, namely the doubly-stochastic matrices.

5. Finding extreme points. Our goal in this section is to explain how the characterization of the convex hull of exact orthogonal solutions from Theorem 2 can be used to obtain *all* exact solutions.

The main observation is that according to Theorem 2, the convex hull of the exact orthogonal solutions is the intersection of the feasible set of the relaxation with the hyperplane defined by constraining the objective to be zero. Therefore, it is possible to preform convex optimization over this set. More specifically, we randomly draw a matrix $W \in \mathbb{R}^{d^2}$ with a uniform distribution over the unit sphere (w.r.t. the Frobenius norm) and obtain exact orthogonal solutions by solving the following optimization problem:

$$\max_{\mu} \quad \mathrm{tr}\, W\mu[R] \tag{12a}$$

$$s.t. \quad \mu \in \mathcal{M}_F \tag{12b}$$

$$\mu[\|RP - QX\|_F^2] = 0 \tag{12c}$$

We now show that the random algorithm returns a unique exact solution with probability one, and the probability of obtaining each exact solution is uniform.

Proof of Theorem 3. Let $(R(i), X(i)), i = 0, \ldots, r-1$ be the members of $\mathrm{ISO}(P,Q)$. We begin by considering the case $P = Q$. In this case $\mathrm{ISO}_R(P,P)$ is a group and we index the group so that $R(0) = I_d$.

A maximizer μ for (12) satisfies $\mu[R] = R(i)$ iff W is a member of the set

$$\mathcal{A}_i = \{W | \,\mathrm{tr}(W^TR(i)) > \mathrm{tr}(W^TR(j)), \forall j \neq i\}$$

clearly the union of $\mathcal{A}_i$ is a disjoint union, and has probability one. We note that $R(\ell)\mathcal{A}_0 = \mathcal{A}_\ell$. Additionally, the probability of the sets $\mathcal{A}_\ell$ is preserved under linear isometries of the space $\mathbb{R}^{d\times d}$ endowed with the Frobenius inner product, and the map $W \mapsto R(\ell)W$ is such an isometry. Therefore the sets $\mathcal{A}_\ell = R(\ell)\mathcal{A}_0$ all have the same probability.

The general case where $P \neq Q$ follows from the fact that $R(0)^T \mathrm{ISO}_R(P,Q) = \mathrm{ISO}_R(P,P)$.

To conclude the argument, note that by lemma 8 $\mu[R] = R(i)$ implies that $\mu[X] = X(i)$. ∎

5.1. Stability. In applications PM problems are usually contaminated with a certain amount of noise. It is therefore important to verify that the solutions of near-exact problems can be recovered as well. To this end, we now consider the case of point clouds $P(\delta), Q(\delta)$ which are obtained by perturbing point clouds $P(0) = P$ and $Q(0) = Q$ which fulfill the conditions of Theorem 2.

We propose a modified version of the algorithm described above to obtain all solutions of the unperturbed problem. Note that generally $PM(P(\delta), Q(\delta))$ won't have exact convex solutions (solutions with zero optimal objective value), so that (12) will become infeasible due to (12c). Therefore we relax this constraint and instead require that the objective of PM-SDP is smaller than some suitable $\epsilon(\delta)$. This results in the following optimization problem:

$$\max_{\mu} \quad \mathrm{tr} W \mu[R] \tag{13a}$$

$$s.t. \quad \mu \in \mathcal{M}_F \tag{13b}$$

$$\mu[\|RP(\delta) - Q(\delta)X\|_F^2] \leq \epsilon(\delta) \tag{13c}$$

An optimal choice for $\epsilon(\delta)$ is

$$\bar{\epsilon}(\delta) = \max_{i=0,\ldots,r-1} \|R(i)P(\delta) - Q(\delta)X(i)\|_F^2$$

This choice assures that all members of $\mathrm{ISO}(P,Q)$ are in the feasible set, while excluding as many irrelevant solutions as possible.

We claim that when $P(\delta), Q(\delta)$ are close to P, Q, we obtain close to exact solutions, with close to uniform probability. To make this statement more precise we need to introduce some notation. We denote the minimizers of (13) by $\mu_\delta(W)$, and their R coordinate by $R_\delta(W)$. Note that both $\mu_\delta(W)$ and $R_\delta(W)$ are sets which may generally contain more than one solution. Also note that when $\delta = 0$ the optimization problems (13) and (12) are identical, and so $R_0(W)$ is uniquely defined almost everywhere and attains all exact orthogonal solutions with uniform probability due to theorem 3. The solutions $R_\delta(W)$ are close to the unperturbed solution $R_0(W)$ if

$$r_\delta(W) = \inf\{r| \quad R_\delta(W) \subseteq B_r(R_0(W))\}$$

is small. We would like to show that for any $\eta > 0$,

$$\mathrm{Prob}\{W| \, r_\delta(W) \geq \eta\} \xrightarrow{\delta \to 0} 0$$

In other words, we claim that r_δ converges in probability to zero. We prove a slightly stronger claim

PROPOSITION 9. *Assume* $\epsilon(\delta) \geq \bar{\epsilon}(\delta)$ *and* $\epsilon(\delta) \xrightarrow{\delta \to 0} 0$. *Then*

$$r_\delta \xrightarrow{a.s.} 0$$

In Appendix A we show that r_δ is Borel measurable and so the probabilistic language used above is justified. The proposition above is essentially a known result on perturbations of SDPs (see [7] pages 492-493). We include the proof for completeness.

Proof. We show pointwise convergence for every W such that $R_0(W)$ is uniquely defined. Assume by contradiction that there is a sequence $\delta_n \longrightarrow 0$ such that $r_{\delta_n}(W) \longrightarrow a > 0$. We can then choose $\mu_n \in \mu_{\delta_n}(W)$ such that

$$|\mu_n[R] - R_0(W)| \longrightarrow a \tag{14}$$

By moving to a subsequence, we can assume that μ_n converges to some μ. We obtain a contradiction to (14) by showing that $\mu[R] = R_0(W)$. This is because

$$\operatorname{tr} W R_0(W) \overset{(*)}{\leq} \operatorname{tr} W \mu_n[R] \longrightarrow \operatorname{tr} W \mu[R] \overset{(**)}{\leq} \operatorname{tr} W R_0(W)$$

where $(*)$ follows from the fact that by assumption $R_0(W)$ is a feasible point of (13), and $(**)$ follows from the fact that μ is a feasible point of (12) since $\epsilon(\delta) \longrightarrow 0$. It follows that μ is a maximizer of (12) and therefore $\mu[R] = R_0(W)$ due to the uniqueness of $R_0(W)$. ∎

6. Experiments. In this section we present experiments illustrating our theoretical results, and experimentally explore the exactness of PM-SDP in cases not covered by our analysis, *i.e.*, when either the simple spectrum condition or the faithfulness condition fail.

To illustrate our results for the asymmetric case we conducted the following experiment. A point cloud P was generated from an i.i.d. normal distribution, and then centralized. Q was obtained by applying random R, X to P. We then solved PM-SDP on the generated point cloud. This process was repeated one-thousand times and in all experiments the unique exact solution with objective value 0 was obtained as can be seen in Figure 4(a).

To illustrate our results for the symmetric case we conducted an experiment which we have already mentioned earlier (Figure 2). We solve PM-SDP on planar point clouds $P = Q$ whose symmetries are

$$R_0 = I_2, \quad R_1 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$$

As predicted, the R coordinate of PM-SDP is in the convex hull of $R_i, i = 0, 1$. We then ran one-hundred iterations of the random algorithm described in Section 5. As predicted, in all iterations one of the solutions R_i was attained (up to a small numerical error). We also proved that R_0 and R_1 should be attained with uniform probability. In this case R_0 was attained in 51 experiments and R_1 in 49 experiments.

We now discuss the cases which are not covered by our analysis. We begin by applying PM-SDP to the asymmetric unfaithful example from Section 1. The results are shown on the top row of Figure 3. We solve PM-SDP, taking P (green) to be the unfaithful shape, and Q (blue) to be a relabeled, rotated version of P. It turns out that although P doesn't meet the conditions of Theorem 2 due to unfaithfulness, PM-SDP is still able to recover the correct orthogonal transformation and permutation. Indeed, as can be seen on the right, applying the obtained transformation to P we are able to align the two shapes perfectly.

In contrast, we provide an example for a shape with multiple eigenvalues where exact recovery is not obtained. These results are shown in the bottom row of Figure 3. The input shape P is asymmetric, but applying PM-SDP to P (green) and Q (blue) chosen to be a rotated, relabeled version of P, does not yield a solution in G. To obtain a solution in G we used the projection procedure described in [14]. The obtained

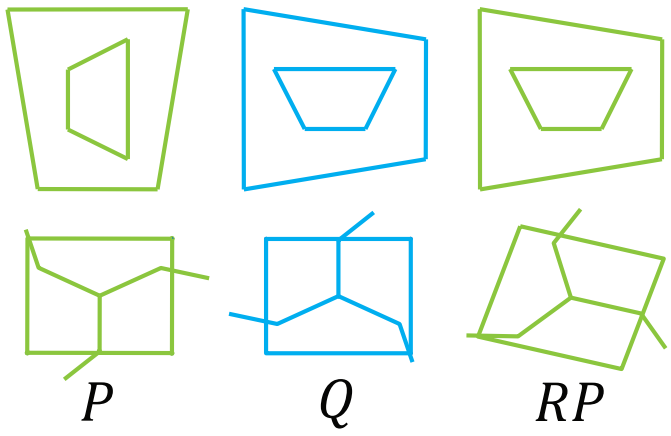

FIG. 3. *Exact recovery when the conditions of Theorem 2 aren't met. By applying PM-SDP to the unfaithful point clouds P (green, top left) and Q (top blue) the correct orthogonal transformation and permutation taking P to Q can be reconstructed (RP, green, top right). The bottom row shows that PM-SDP fails to align the shapes P and Q which violate the simple spectrum condition.*

solution is still not the correct exact solution as shown in the image on the bottom right which is not aligned correctly.

We also conducted some random experiments. The results are summarized in Figure 4. Each experiment was run one-thousand times on point clouds P and Q obtained from P by applying a random orthogonal transformation and permutation.

Experiment (b) quantitatively shows that exact recovery is obtained for unfaithful shapes as well. Each P here is generated by superimposing two shapes with different reflectional symmetry as in the unfaithful example in Figure 3.

Experiment (c) quantitatively shows that exact recovery isn't always obtained for shapes with eigenvalue multiplicity. Here the two dimensional asymmetric shapes P are generated by superimposing two shapes with non-bilateral symmetry groups. In all cases PM-SDP returns a solution whose R, X coordinates aren't in G, and we project onto G using the scheme described in [14]. It can be seen that in some cases the PM objective obtained from R, X after projection is non-zero although an exact solution does exist.

Experiment (d) examines a different class of shapes with eigenvalue multiplicity for which exact recovery is obtained. Here P is generated from an i.i.d. normal distribution. It is then centralized and scaled along one of its principal axes so that a shape with eigenvalue multiplicity is obtained. In all cases PM-SDP returned the correct exact solution of the problem (even before projection).

To summarize, we are not aware of simple spectrum, unfaithful shapes, for which PM-SDP does not achieve exact recovery. Thus a possible direction for future work is investigating whether this condition can be removed or weakened. In contrast we provide examples of shapes with double eigenvalue multiplicity for which exact recovery fails, and thus the simple spectrum condition is tight.

7. Exact graph matching. The purpose of this section is establishing that PM, and in fact even the subproblem of exact PM, is computationally hard. We do this by showing equivalence of exact PM and exact graph matching.

For our discussion here we define the exact PM problem as the problem of deciding whether $\mathrm{d}(P, Q) = 0$ for given point clouds $P, Q \in \mathbb{R}^{d\times n}$. The input of the exact graph

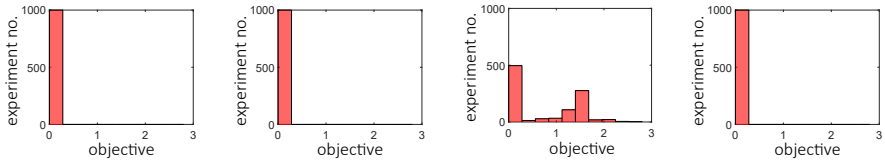

(a) (b) (c) (d)

FIG. 4. *Exact recovery experimental evaluation. When the conditions of the (asymmetric) exact recovery theorem are met, an optimal solution with objective zero is always attained (a). Exact recovery is also obtained for randomly generated unfaithful point clouds (b). For randomly generated shapes with eigenvalue multiplicity exact recovery can fail (c). In (d) successful exact recovery is shown for an easier class of shapes with eigenvalue multiplicity.*

matching (GM) problem is a pair of graphs $G_A = (V_A, A)$ and $G_B = (V_B, B)$. The vertices of the graphs are assumed to be of equal cardinality n. The matrices A, B are assumed to be symmetric and are chosen to reflect relevant information on the graph. The goal of the exact graph matching problem is deciding whether there is a relabeling of the vertices so that the graphs are identical. Mathematically this amounts to deciding whether there is a permutation matrix $X \in \Pi_n$ such that

$$XA = BX \tag{15}$$

In the special case where A, B are the adjacency matrices of the graphs, this problem is called the graph isomorphism problem. The graph isomorphism problem is in NP, but is not known to be either NP-complete or in P. We now prove the equivalence of exact PM and exact GM:

proof of Theorem 1. The main observation needed for the proof is that if $P, Q \in \mathbb{R}^{d\times n}$ satisfy $P^TP = A$ and $Q^TQ = B$, then the set of solutions of (15) is exactly the projection onto the X coordinate of the set of solutions to the exact PM equation

$$RP = QX \tag{16}$$

To see this note that if (16) holds then

$$XA = X(RP)^TRP = X(QX)^TQX = BX$$

On the other hand, if (15) holds then

$$P^TP = X^TQ^TQX$$

and so P and QX are two factorizations of the same matrix. Since factorization is unique up to an orthogonal transformation there must be some $R \in \mathcal{O}(d)$ such that (16) holds.

We now prove the theorem using this observation. We begin by showing a polynomial reduction from exact GM to exact PM. Assume we are given two graphs defined by the matrices A, B. We note that we can assume w.l.o.g. that $A, B \succeq 0$ since for any $\lambda \in \mathbb{R}$, a permutation X satisfies (15) if and only if

$$X(A + \lambda I_n) = (B + \lambda I_n)X$$

We can therefore factorize A, B to obtain P, Q, and the observation implies that the graphs are isomorphic iff $d(P, Q) = 0$.

For the opposite reduction- assume we are given point clouds P, Q. We define $A = P^TP$ and $B = Q^TQ$. Our observation implies that $d(P, Q) = 0$ if and only if there is a solution to (15). ∎

Appendix A. Measurability. We show r_δ is upper semi-continuous at almost every W and hence Borel measurable. Assume by contradiction that there is some sequence $W_n \longrightarrow W$ such that $r_\delta(W_n) \longrightarrow a > r_\delta(W)$. We can then choose $R_n \in R_\delta(W_n)$ such that

$$|R_n - R_0(W)| \longrightarrow a \tag{17}$$

by moving to a subsequence we can assume that R_n converges to some R, using similar reasoning as in the proof of Theorem 9 it can be shown that $R \in R_\delta(W)$. It follows that $|R - R_0(W)| \leq r_\delta(W)$ in contradiction to (17).

REFERENCES

[1] Y. Aflalo, A. Bronstein, and R. Kimmel, *On convex relaxation of graph isomorphism*, Proceedings of the National Academy of Sciences, 112 (2015), pp. 2942–2947.

[2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, *Scape: shape completion and animation of people*, in ACM Transactions on Graphics (TOG), vol. 24, ACM, 2005, pp. 408–416.

[3] L. Babai, *Graph isomorphism in quasipolynomial time*, CoRR, abs/1512.03547 (2015).

[4] L. Babai, D. Y. Grigoryev, and D. M. Mount, *Isomorphism of graphs with bounded eigenvalue multiplicity*, in Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, STOC '82, 1982, pp. 310–324.

[5] A. I. Barvinok, *Problems of distance geometry and convex properties of quadratic maps*, Discrete & Computational Geometry, 13 (1995), pp. 189–202.

[6] P. J. Besl and N. D. McKay, *Method for registration of 3-d shapes*, in Robotics-DL tentative, International Society for Optics and Photonics, 1992, pp. 586–606.

[7] J. F. Bonnans and A. Shapiro, *Perturbation analysis of optimization problems*, Springer Science & Business Media, 2013.

[8] M. A. Fischler and R. C. Bolles, *Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography*, Communications of the ACM, 24 (1981), pp. 381–395.

[9] Q.-X. Huang and L. Guibas, *Consistent shape maps via semidefinite programming*, in Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing, SGP '13, 2013, pp. 177–186.

[10] I. Kezurer, S. Z. Kovalsky, R. Basri, and Y. Lipman, *Tight relaxation of quadratic matching*, Comput. Graph. Forum, 34 (2015), pp. 115–128.

[11] Y. Khoo and A. Kapoor, *Towards non-iterative closest point: Exact recovery of pose for rigid 2d/3d registration using semidefinite programming*, CoRR, abs/1501.00630 (2015).

[12] J. B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM J. on Optimization, 11 (2000), pp. 796–817.

[13] Z.-Q. Luo, W.-k. Ma, A. M.-C. So, Y. Ye, and S. Zhang, *Semidefinite relaxation of quadratic optimization problems*, Signal Processing Magazine, IEEE, 27 (2010), pp. 20–34.

[14] H. Maron, N. Dym, I. Kezurer, S. Z. Kovalsky, and Y. Lipman, *Point registration via efficient convex relaxation*, To appear in ACM Siggraph, (2016).

[15] M. Ovsjanikov, J. Sun, and L. Guibas, *Global intrinsic symmetries of shapes*, in Computer graphics forum, vol. 27, Wiley Online Library, 2008, pp. 1341–1348.

[16] S. Rusinkiewicz and M. Levoy, *Efficient variants of the ICP algorithm*, Proceedings Third International Conference on 3-D Digital Imaging and Modeling, (2001), pp. 145–152.

[17] A. M.-C. So and Y. Ye, *Theory of semidefinite programming for sensor network localization*, Mathematical Programming, 109 (2007), pp. 367–384.